

# Plato HPC cluster

Plato is a Linux-based, heterogeneous, high-performance computing (HPC) cluster at the University of Saskatchewan. It is used for research projects and for training, and is restricted to USask users and their collaborators. Plato is managed by ICT's Advanced Research Computing (ARC) team.

Although it is not a Compute Canada cluster, Plato is configured to be as similar as possible to Compute Canada clusters. For instance, Plato has the same scientific and general software stack, and uses the same scheduler.

Plato cannot provide all the computing power required for USask research projects. Researchers with important computing needs should consider Plato a stepping stone to test their projects locally before moving to Compute Canada clusters.

## Quick links

- Prospective users should check [Getting access to Plato](#).
- Newcomers should read [Getting started on Plato](#).
- [Subpages](#) contain all documentation not on this overview page.
- Most [Compute Canada Documentation](#) also applies to Plato due to system similarities.

## Contents

- [Basic information](#)
- [Storage](#)
- [Network](#)
- [Scheduler](#)
- [Nodes](#)
  - [Choosing nodes](#)
  - [Login nodes](#)
- [Educational accounts](#)
- [Other topics \(subpages\)](#)

## Basic information

|                      |   |
|----------------------|---|
| <b>SSH hostname</b>  | plato.usask.ca (VPN required when off-campus)                         |
| <b>Globus access</b> | usask#gateway /plato/<nsid> (home directory)                          |
| <b>System status</b> | <a href="#">ARC main page</a> , and <code>cluster-info</code> command |
| <b>Storage usage</b> | <code>quota</code> command  |

## Storage

|   |  |
|---|--|
| <b>Home filesystem</b><br><code>/home</code>        | <ul style="list-style-type: none"><li>▪ 24T total volume</li><li>▪ Location of home directories</li><li>▪ Each home directory has a small (300G) quota</li><li>▪ No backup</li><li>▪ For active research data and user programs</li><li>▪ Inactive research data must be moved to Datastore</li></ul>  |
| <b>Datastore service</b><br><code>/datastore</code> | <ul style="list-style-type: none"><li>▪ 723T total volume</li><li>▪ Also accessible from outside Plato</li><li>▪ Faculty members are entitled to 3T free space and can obtain more</li><li>▪ Backed-up</li><li>▪ For long-term storage and inactive research data</li><li>▪ Read-only access on the compute nodes, read-write on the login nodes</li></ul> |
| <b>Compute nodes</b><br><code>/local</code>         | <ul style="list-style-type: none"><li>▪ Temporary storage for jobs</li><li>▪ Total volume varies by compute node type</li><li>▪ Use <code>\$SLURM_TMPDIR</code> to get the temporary directory for your current job, in the form: <code>/local/\$USER/\$SLURM_JOB_ID</code></li><li>▪ Files are deleted when a job ends</li></ul>                          |

## Network

Most Plato compute nodes are interconnected by a 1Gb Ethernet link. Login nodes are connected to compute nodes and to the University network by a 10Gb Ethernet link. Some compute nodes (see below) are interconnected by 10Gb Ethernet or FDR InfiniBand. Both login and compute nodes can access the external network, including the Internet.

## Scheduler

Plato uses the SLURM scheduler. Job duration is limited to 21 days, except GPU (7 days) and large-memory (30 days) nodes. (This limit does not apply to contributed hardware.) Shorter jobs get increased priority, according to the categories below. The default allocation is 20 minutes for a single task on one CPU with 512M of memory. Users can have a maximum of 1000 jobs at a time in the scheduler.

| Maximum duration (d-hh:mm) | Priority factor |
|----------------------------|-----------------|
| 21-00:00                   | 1               |
| 4-00:00                    | 2               |
| 12:00                      | 4               |
| 04:00                      | 8               |

## Nodes

| Count | Type         | Public | Cores per node | Memory per node | CPU  | Highest SIMD | GPU            | /local storage | Interconnect   |
|-------|--------------|--------|----------------|-----------------|--|--------------|----------------|----------------|----------------|
| 64    | "Pipit"      | X      | 16             | 31000M (30G)    | 2 x Intel Xeon E5-2640 v2 @ 2.00GHz "Ivy Bridge" | AVX          | -              | 347G           | 1Gb Ethernet   |
| 32    | "Penguin"    | X      | 40             | 190000M (185G)  | 2 x Intel Xeon Gold 6148 @ 2.40GHz "Skylake"     | AVX512       | -              | 781G           | 10Gb Ethernet  |
| 1     | Large-memory | X      | 48             | 2048000 (2000G) | 4 x Intel Xeon E7-4850 v2 @ 2.30GHz "Ivy Bridge" | AVX          | -              | 3.2T           | 10Gb Ethernet  |
| 2     | GPU          | X      | 16             | 31000M (30G)    | 2 x Intel Xeon E5-2640 v3 @ 2.60GHz "Haswell"    | AVX2         | 2 x NVIDIA K40 | 805G           | 1Gb Ethernet   |
| 20    | GWF          | -      | 16             | 31000M (30G)    | 2 x Intel Xeon E5-2640 v2 @ 2.00GHz "Ivy Bridge" | AVX          | -              | 347G           | 1Gb Ethernet   |
| 2     | Tse group    | -      | 32             | 250000M (244G)  | 2 x Intel Xeon E5-2683 v4 @ 2.10GHz "Broadwell"  | AVX2         | 2 x NVIDIA K80 | 768G           | FDR InfiniBand |

## Choosing nodes

Plato will choose the appropriate node type for your job according to your resource requirements (cores per node, memory, GPUs). It is therefore not necessary to request a specific node type. Doing so may reduce the number of nodes eligible to run your job, increasing your wait time. You can specify a node type using SLURM options: `--constraint=ivybridge` (for Pipit nodes) or `--constraint=skylake` (for Penguin nodes).

## Login nodes

Plato has a single login node, `plato1gn01`, accessible by SSH at `plato.usask.ca`. The login node should be used to prepare jobs and submit them to the scheduler, to compile programs, and to run short calculations that require little memory and processing power. Intensive processes must never be run on the login node; they must be submitted to the scheduler to run on the compute nodes. Each user has access to only 8 CPU cores and 16G of memory on the login node to ensure it remains responsive for all users.

## Educational accounts

If you are granted access to Plato for a class or training workshop, your jobs will be limited to 12 hours of runtime, and you will not have access to the large-memory nodes.

Other topics (subpages)